1 5G NTN (Satellite) In-Flight Connectivity for Aircraft: Impact of LEO Constellation Size on Throughput, SINR, and Latency

1.1 Overview

In-flight connectivity over a 5G non-terrestrial network is delivered by satellites that are always in motion. Constellation size is the basic dimensioning question: with too few satellites the aircraft spends part of the flight with no serving satellite, and every added satellite changes coverage, handover frequency, and delivered performance. This case study examines that question in NetSim. An aircraft flies from Bangalore to Delhi during a 2.20-hour simulation while the LEO constellation moves along its orbits, and the serving satellite changes repeatedly on the way. Two runs differ in a single parameter: 50 satellites against 100, on the same Ka-band configuration. The study follows the serving-satellite timeline and the throughput, downlink SINR, and latency seen by the aircraft, then quantifies the effect of doubling the constellation: more of the flight with service, higher delivered throughput, and lower delay.

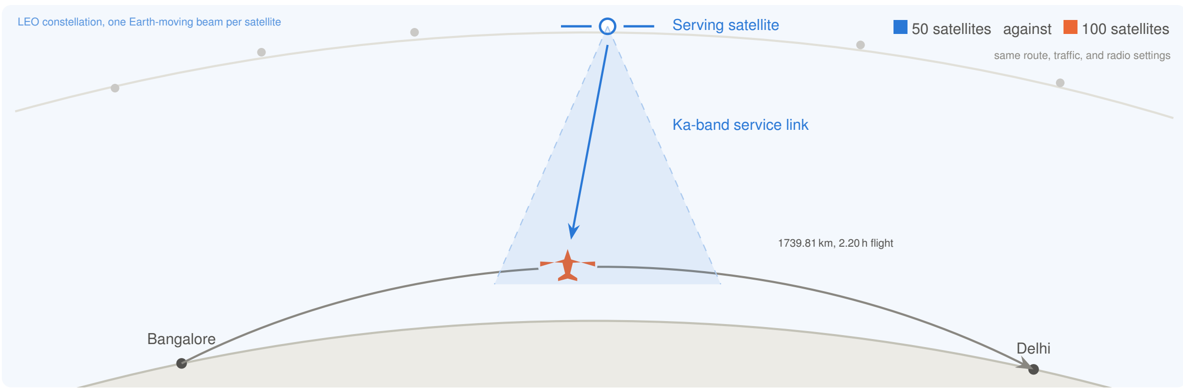

Figure 1-1: *The scenario at a glance. An aircraft flies the 1739.81 km Bangalore to Delhi route in 2.20 hours while a Ka-band LEO constellation passes overhead. At any instant one satellite serves the aircraft through its Earth-moving beam. The study repeats the same flight with 50 and with 100 satellites and compares the service the aircraft receives.*

1.2 Problem statement

Service to an aircraft in a non-terrestrial network depends on satellite visibility, elevation angle, beam geometry, and radio-link quality. Over a long route every one of these varies, so the satellite serving the aircraft changes during the flight. The objective of this study is to track that satellite association over time and to evaluate the radio and application performance the aircraft sees, first with 50 satellites and then with 100.

1.3 Approach

The aircraft follows a Bangalore to Delhi flight path. Traffic is downlink Full Buffer over UDP, so the aircraft receives as much downlink traffic as the link can deliver whenever it is connected. Both runs use Ka-band NR NTN with a Gaussian satellite antenna model, free-space path loss, and a 10 degree elevation mask. Each satellite carries one Earth-moving beam (beam type `Earth_Moving` in NetSim) that travels with it, so coverage at any point on the route depends on where the satellites are at that moment. The two runs are identical except for constellation size. The analysis compares their throughput, downlink SINR, and latency.

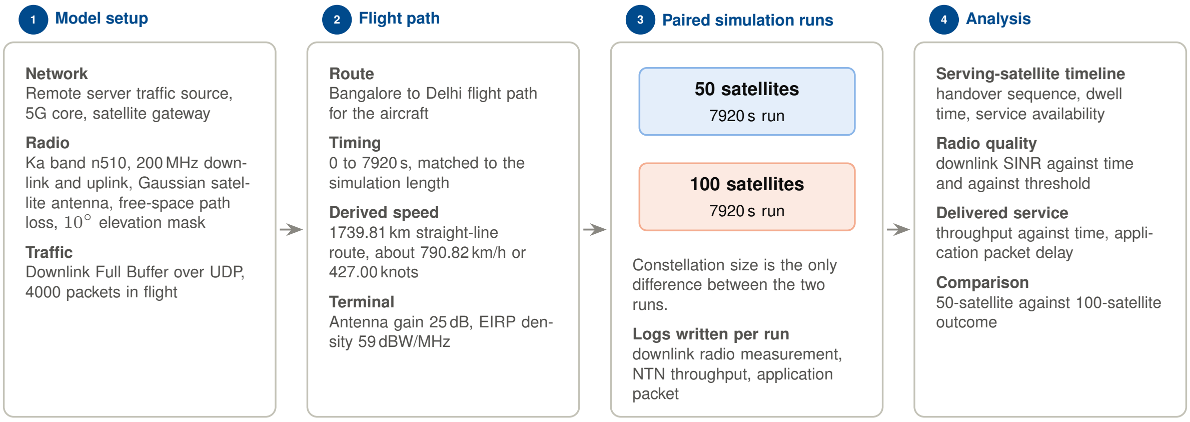

Figure 1-2: *Study workflow. Model setup and the flight path are configured once and held fixed. Constellation size is the only parameter that differs between the two runs, so the difference in the results isolates the effect of satellite density on aircraft service.*

1.4 Simulation setup

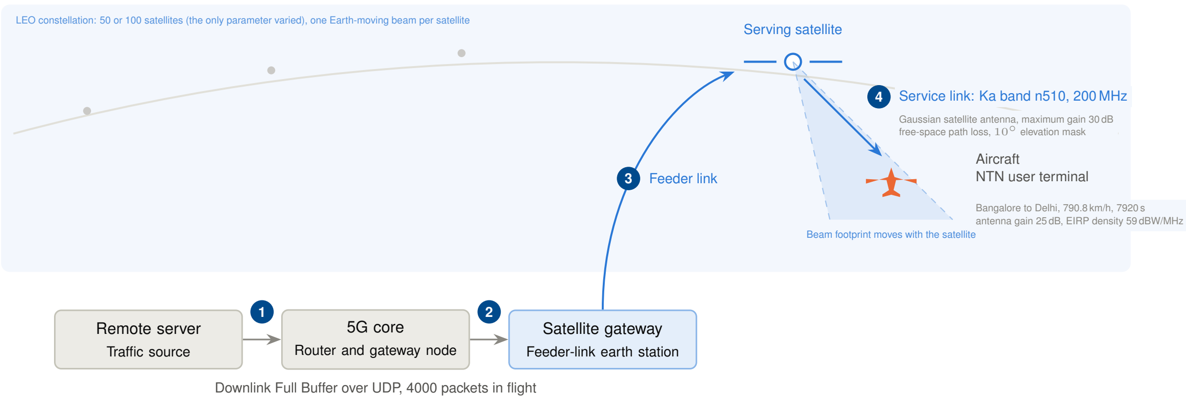

Figure 1-3: *End-to-end downlink path, numbered in traffic order. Full Buffer traffic originates at the remote server (Wired_Node_9 in NetSim) (1), crosses the 5G core to the satellite gateway (Satellite_Gateway_11) (2), climbs the feeder link to the serving satellite (3), and reaches the aircraft (Aircraft_10) over the Ka-band service link (4). Each satellite carries one Earth-moving beam, so the footprint travels with the satellite and the aircraft is served only while it lies inside the beam.*

Table 1-1: *Scenario configuration*

Parameter	Value
Aircraft	Aircraft_10
Mobility model	File Based Mobility, time-stamped waypoints for Aircraft_10
Start point	Bangalore
End point	Delhi
Simulation time	7920 s
Flight time window	0–7920 s

Parameter	Value
Flight route	Direct route from Bangalore to Delhi
Route speed	About 790.82 km/h, or 427.00 knots
Traffic	Downlink Full Buffer from the remote server (Wired_Node_9 in NetSim) to Aircraft_10 using UDP, packets in flight 4000
Radio band and bandwidth	Ka band, n510, 200 MHz downlink and 200 MHz uplink
Satellite antenna model	Gaussian antenna model, beam type Earth_Moving
Terminal and link budget	Aircraft antenna gain 25 dB, EIRP density 59 dBW/MHz, maximum antenna gain 30 dB
Satellite path loss model	Free_space with elevation mask 10 degrees
Updated cases	50-satellite and 100-satellite constellations with the same flight path and traffic settings

1.5 Network topology and flight path

The aircraft follows the configured flight path, defined through File Based Mobility. Each waypoint gives a simulation time and an aircraft position, and NetSim moves the aircraft to that position at that instant. The route leaves Bangalore at 0 s and reaches Delhi at 7920 s, the full simulation duration, so the aircraft is in motion for the entire 2.20-hour run.

No fixed speed is configured. NetSim derives speed from the distance between consecutive waypoints and the time gap between them. The straight-line route length is about 1739.81 km, which gives an average speed of about 790.82 km/h, or 427.00 knots.

1.6 Why the serving satellite changes

The serving satellite changes many times during the flight. Elevation angle, slant range, and beam geometry all vary as the aircraft and the satellites move along their trajectories, and each of them feeds directly into the received SINR. A satellite that serves the aircraft well at one point on the route sets below the elevation mask a few minutes later, and service passes to whichever satellite then offers a link above the mask.

1.7 Serving-satellite timeline and service availability

Table 1-2 and Table 1-3 list the serving satellite observed for the aircraft over time. Figure 1-7 places both timelines on a common time axis, so the handover sequence, the dwell time on each satellite, and the intervals with no serving satellite can be compared directly.

Two points follow from the timelines. First, the aircraft alternates between orbital planes PL-1 and PL-7 for most of the route: the handover sequence is driven by which plane happens to have a satellite above the mask, rather than by a single plane tracking the aircraft. Second, the denser constellation shortens the individual serving intervals and makes the gaps between them fewer and narrower, so handovers occur more often and coverage is more continuous.

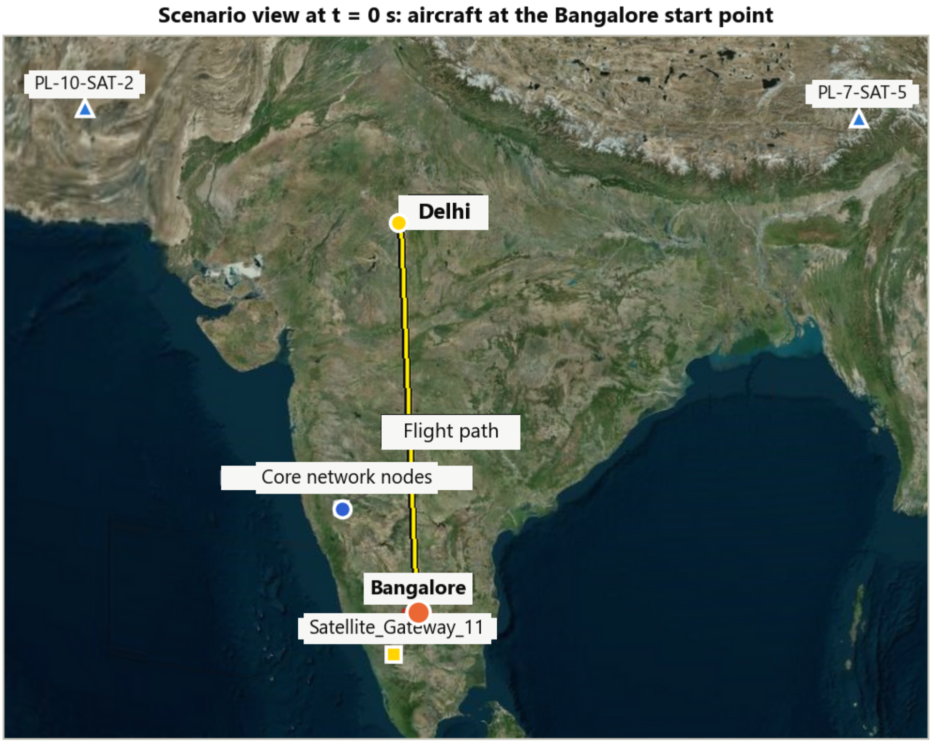

Figure 1-4: *Scenario view at the start of the simulation, with the aircraft at the Bangalore start point at 0 s. The yellow line marks the configured flight path toward Delhi over the 0–7920 s simulation. The same view shows the core network nodes, the satellite gateway, and two of the visible satellites.*

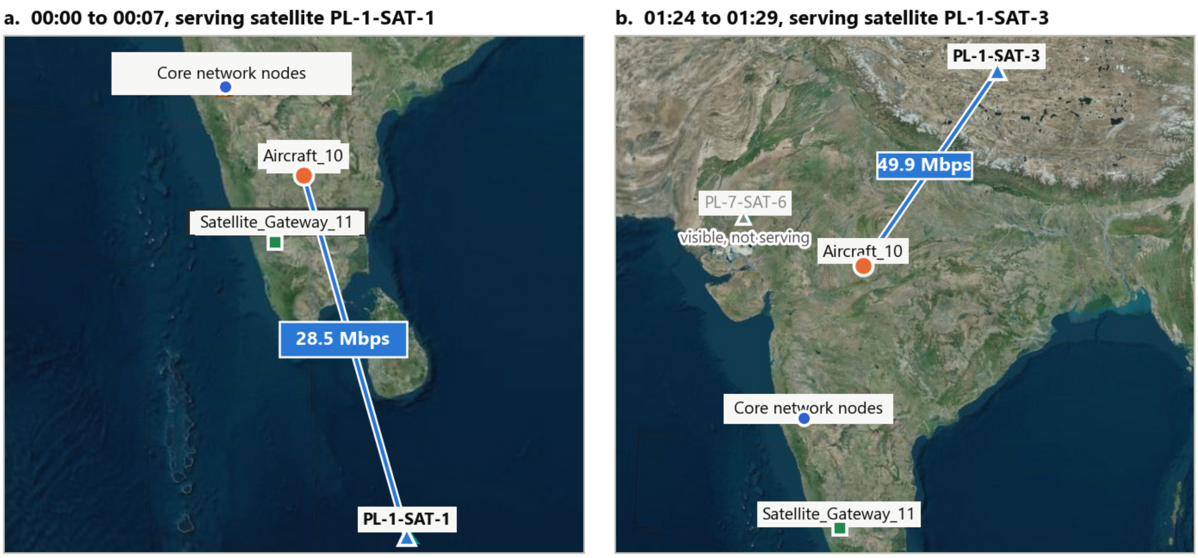

Figure 1-5: *Two snapshots during the flight in the 100-satellite run. The blue link and the throughput badge identify the serving satellite at that instant. Early in the flight the aircraft is served by PL-1-SAT-1 at 28.5 Mbps; about an hour and a half later it is served by PL-1-SAT-3 at 49.9 Mbps, with other satellites visible but not serving. Serving-satellite intervals are taken from Table 1-3.*

Table 1-2: *Serving satellite timeline for Aircraft_10 in the updated 50-satellite case*

Time (h:min)	Serving satellite	Peak SINR (dB)
00:00–00:07	PL-1-SAT-1	-4.33
00:20–00:28	PL-1-SAT-5	-0.53

Time (h:min)	Serving satellite	Peak SINR (dB)
00:37–00:43	PL-7-SAT-1	-3.57
00:44–00:48	PL-1-SAT-4	3.14
00:56–01:03	PL-7-SAT-5	-1.48
01:04–01:09	PL-1-SAT-3	15.31
01:15–01:23	PL-7-SAT-4	30.18
01:24–01:29	PL-1-SAT-2	22.28
01:34–01:42	PL-7-SAT-3	-0.32
01:44–01:50	PL-1-SAT-1	4.00
01:51–01:53	PL-8-SAT-2	-5.87
01:54–02:01	PL-7-SAT-2	-6.23
02:02–02:10	PL-1-SAT-5	0.34

The serving sequence is derived from the downlink radio measurement log. Gaps between table entries are periods with no serving satellite.

Table 1-3: *Serving satellite timeline for Aircraft_10 in the 100-satellite case*

Time (h:min)	Serving satellite	Peak SINR (dB)
00:00–00:07	PL-1-SAT-1	-3.65
00:13–00:17	PL-1-SAT-10	-2.16
00:20–00:28	PL-1-SAT-9	-2.79
00:28–00:33	PL-7-SAT-2	-5.83
00:34–00:38	PL-1-SAT-8	-4.03
00:38–00:43	PL-7-SAT-1	-3.52
00:43–00:48	PL-1-SAT-7	-0.14
00:49–00:53	PL-7-SAT-10	-2.52
00:53–00:58	PL-1-SAT-6	4.68
01:00–01:03	PL-7-SAT-9	-1.22
01:03–01:08	PL-1-SAT-5	3.82
01:08–01:13	PL-7-SAT-8	3.89
01:14–01:18	PL-1-SAT-4	18.80
01:18–01:23	PL-7-SAT-7	30.41
01:24–01:29	PL-1-SAT-3	23.49
01:34–01:38	PL-1-SAT-2	5.06
01:38–01:42	PL-7-SAT-5	0.95
01:56–02:00	PL-1-SAT-10	0.34
02:04–02:10	PL-1-SAT-9	1.73

1.8 50-satellite and 100-satellite comparison

Both runs use Ka-band n510, 200 MHz bandwidth, the Gaussian antenna model, aircraft antenna gain 25 dB, EIRP density 59 dBW/MHz, maximum antenna gain 30 dB, and packets in flight 4000. Throughput comes from the NTN throughput log, latency from the application packet log, and SINR from the radio measurement log. The headline comparison appears first; the time-series figures that follow stack the two runs on a common time axis, and the shaded windows mark intervals where that log records no serving satellite.

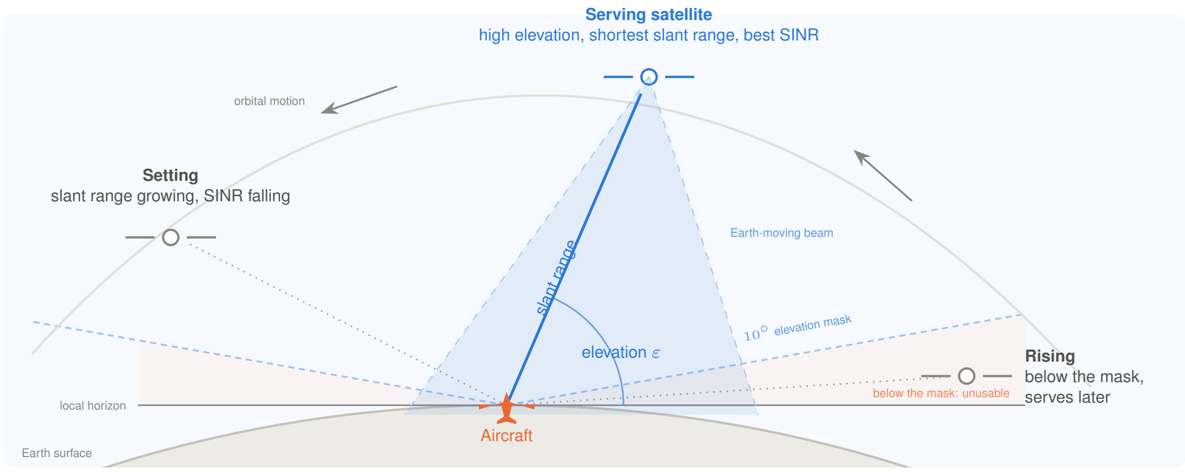

Figure 1-6: *Link geometry seen from the aircraft. A satellite serves the aircraft only while its elevation ε is above the 10 degree mask and the aircraft lies inside its Earth-moving beam. The serving satellite offers the highest elevation and the shortest slant range. The setting satellite loses SINR as its slant range grows, and the rising satellite sits in the shaded below-mask zone, where it cannot serve yet. When no satellite meets both conditions, service is interrupted; these interruptions appear as the shaded windows in the result plots and as the gaps in the serving-satellite timelines.*

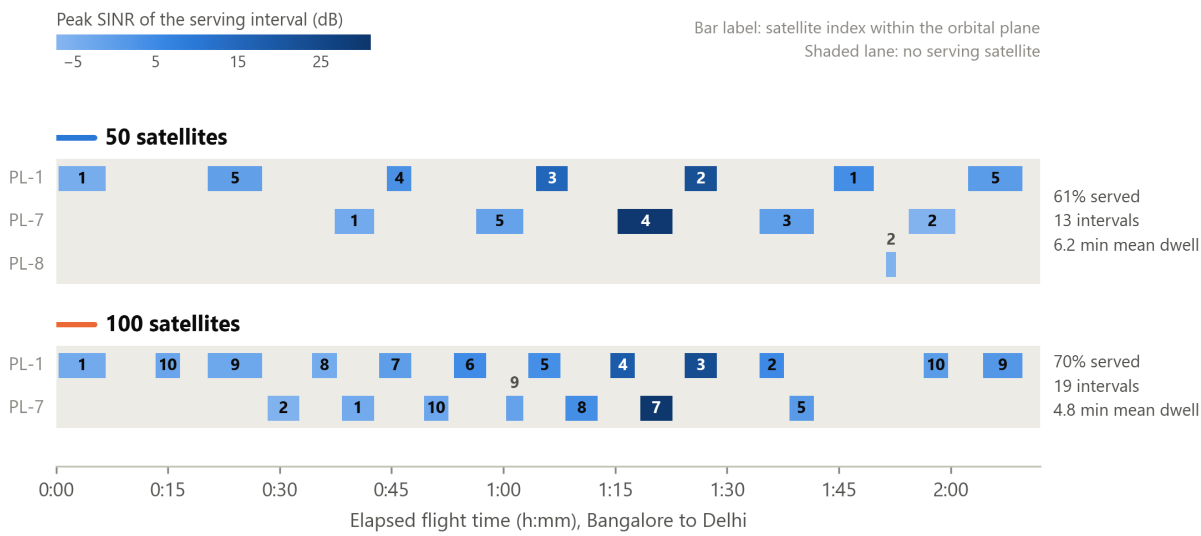

Figure 1-7: *Serving-satellite timeline for both runs, with each interval placed in the lane of its orbital plane and shaded by the peak SINR recorded during that interval. The 50-satellite run holds 13 serving intervals covering 61 per cent of the flight, at 6.2 minutes mean dwell. The 100-satellite run holds 19 shorter intervals covering 70 per cent of the flight, at 4.8 minutes mean dwell: more frequent handovers and less time without a serving satellite.*

Table 1-4: *Throughput summary*

Case	Minimum	Average	Maximum
50 satellites	0.28 Mbps	65.36 Mbps	1033.47 Mbps
100 satellites	0.26 Mbps	73.41 Mbps	1033.47 Mbps

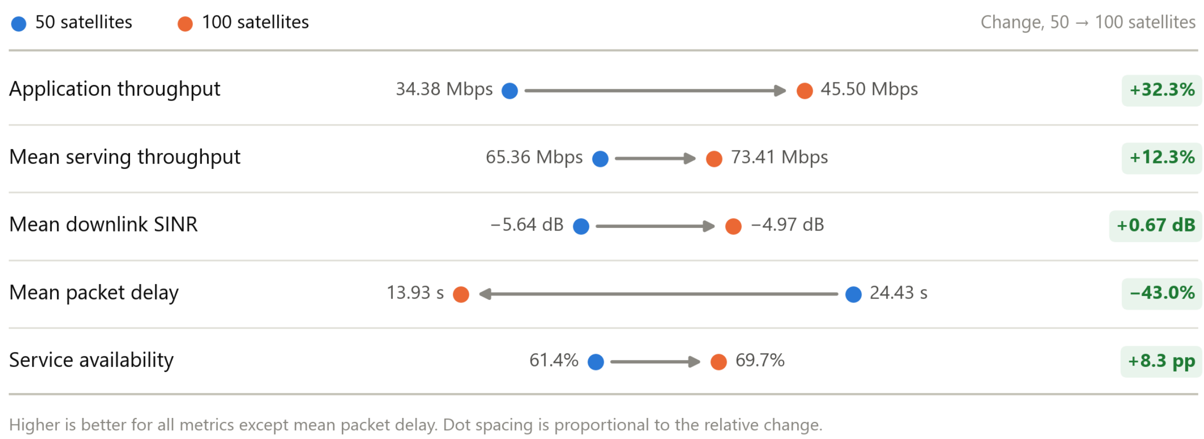

Figure 1-8: *Headline comparison between the 50-satellite and 100-satellite runs. The 100-satellite constellation delivers higher application throughput, higher mean SINR, lower mean packet delay, and a larger share of the flight with a serving satellite. Service availability is the share of the 2.20-hour flight covered by the serving-satellite timelines in Table 1-2 and Table 1-3.*

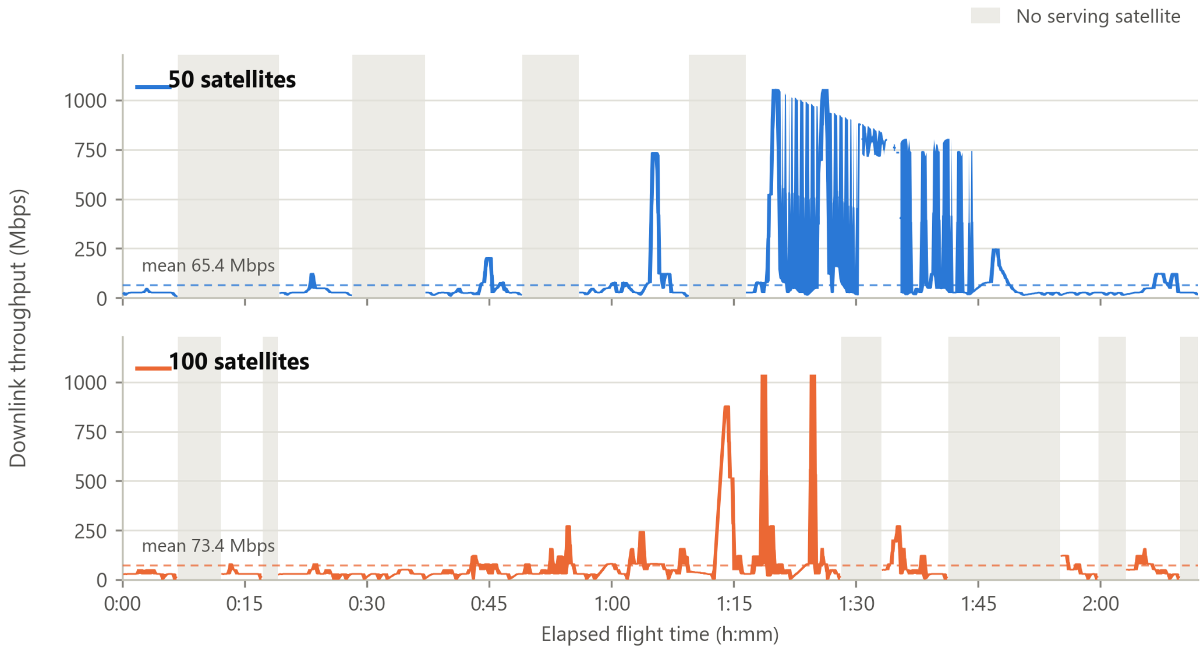

Figure 1-9: *Throughput against time. Both runs sit near 25–50 Mbps for most of the flight and rise to short peaks above 1000 Mbps when the aircraft passes close to beam centre at high elevation. The 50-satellite run reaches its peaks in one long, sustained burst around 1:15–1:45; the 100-satellite run spreads smaller peaks more evenly and has fewer and narrower windows with no serving satellite.*

Table 1-5: *Downlink SINR summary*

Case	Minimum	Average	Maximum
50 satellites	-12.16 dB	-5.64 dB	30.09 dB
100 satellites	-12.24 dB	-4.97 dB	30.41 dB

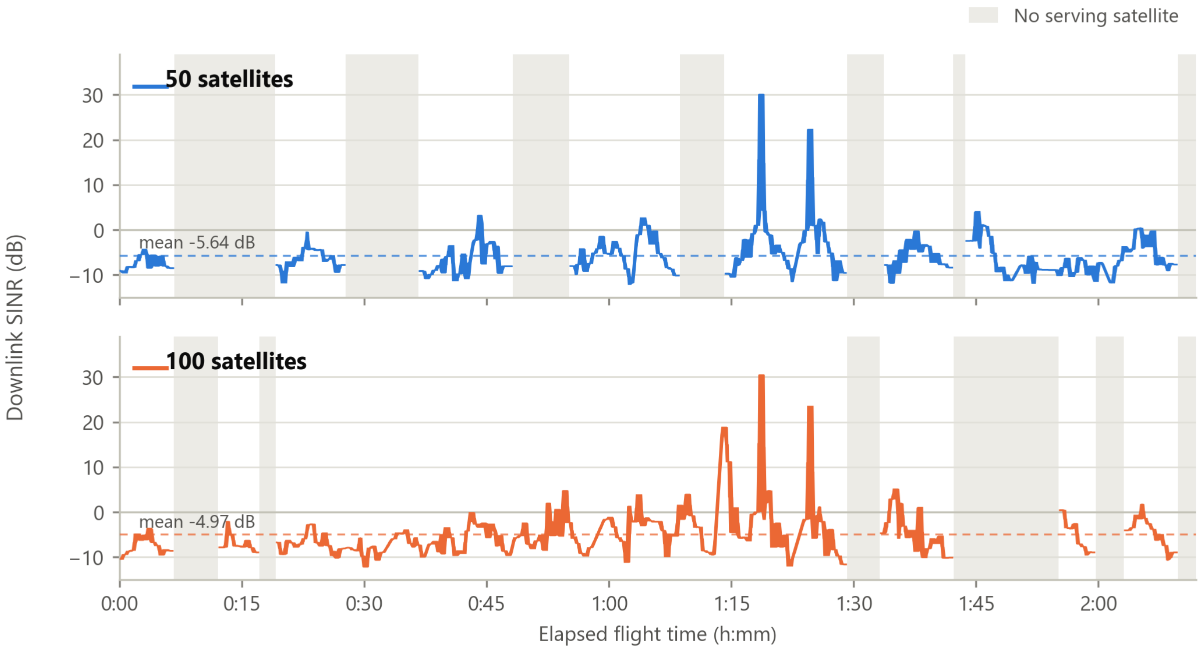

Figure 1-10: *Downlink SINR against time. Both runs spend most of the flight below 0 dB, with the highest SINR occurring in the same 1:15–1:30 window in which the throughput peaks appear. The 100-satellite run holds a slightly higher mean, -4.97 dB against -5.64 dB, and shows the more frequent SINR steps that follow from its shorter serving intervals.*

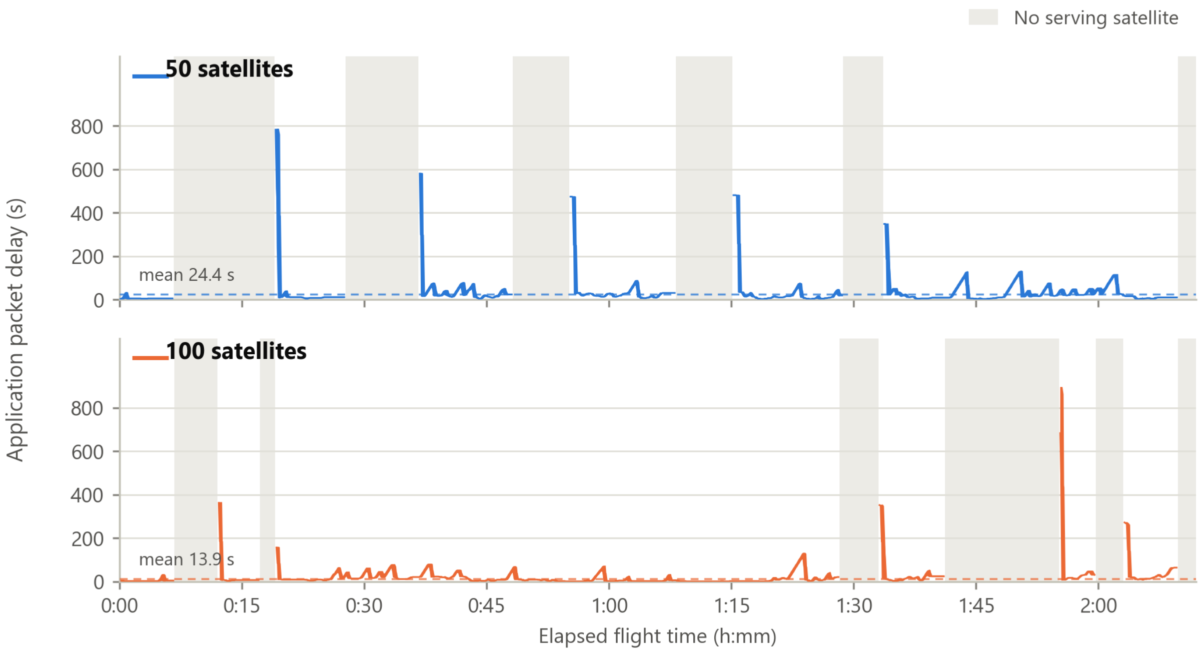

Figure 1-11: *Application packet delay against time. The baseline delay is tens of seconds in both runs. Every large spike occurs immediately after a window with no serving satellite, which identifies these spikes as the backlog delivered once the link is reacquired rather than as propagation delay. The 100-satellite run halves the mean delay, from 24.4 s to 13.9 s.*

Table 1-6: *Latency summary*

Case	Minimum	Average	Maximum
50 satellites	349.55 ms	24429.56 ms	783398.83 ms
100 satellites	6.50 ms	13929.08 ms	893539.12 ms

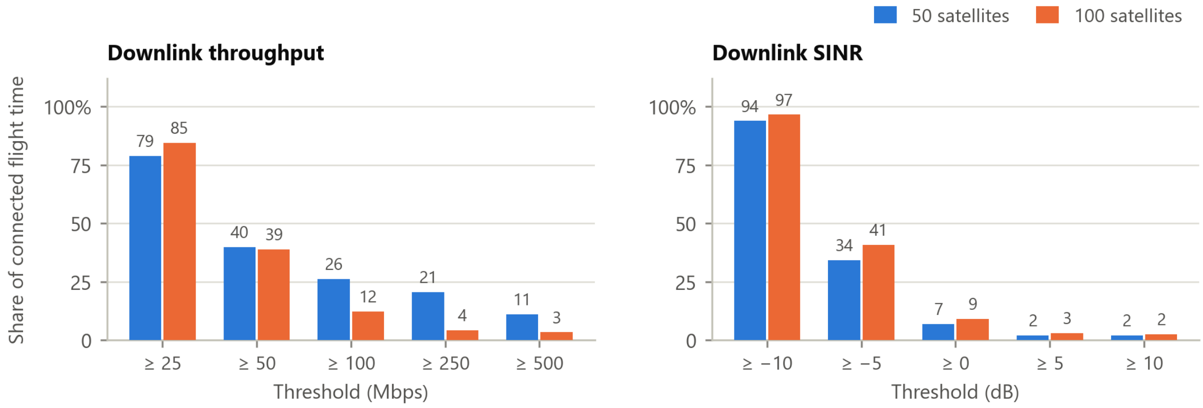

Figure 1-12: *Share of connected flight time above a given throughput and SINR threshold, read from the plotted series. The 100-satellite run is better at the moderate thresholds that determine everyday service: 85 per cent of connected time above 25 Mbps against 79 per cent, and 41 per cent of connected time above -5 dB against 34 per cent. The 50-satellite run holds more time at the high thresholds, because its longer dwell on a single satellite carries it further through the high-elevation part of that satellite's pass.*

Table 1-7: *Application throughput and service availability*

Case	App. throughput	Serving intervals	Mean dwell	Service availability
50 satellites	34.38 Mbps	13	6.2 min	61%
100 satellites	45.50 Mbps	19	4.8 min	70%

The 100-satellite run delivers higher application throughput than the 50-satellite run: 45.50 Mbps against 34.38 Mbps. It receives more packets over the same 7920 s, shows fewer no-throughput windows in the throughput log, and holds a higher mean serving throughput, 73.41 Mbps against 65.36 Mbps.

1.9 Interpretation against aircraft connectivity

Both runs use Ka-band n510, 200 MHz bandwidth, one aircraft, Full Buffer UDP traffic, one Earth-moving beam per satellite, Gaussian antenna gain, and 25 dB aircraft antenna gain. Under this setup, the application throughput is 34.38 Mbps for 50 satellites and 45.50 Mbps for 100 satellites.

The gain from doubling the constellation comes from availability rather than from peak link quality. Peak SINR and peak throughput are almost identical in the two runs: 30.09 dB against 30.41 dB, and 1033.47 Mbps in both. What changes is how much of the flight has a serving satellite at all, 61 per cent against 70 per cent. That extra covered time, together with the shorter interruptions between serving intervals, raises the delivered application throughput and halves the mean packet delay.

These results are in the same range as published aviation broadband figures, though they model one configured scenario rather than any commercial service. Starlink Aviation lists 100–250 Mbps with peaks up to 450 Mbps, Inmarsat Global Xpress lists downlink speeds up to 50 Mbps, and Viasat Ka aviation states speeds typically greater than 30 Mbps. The NetSim value depends on constellation size, Earth-moving beam coverage, Gaussian beamwidth, aircraft antenna gain, path loss, SINR, and packet delivery over the flight duration.

The latency values in this study are application packet delay from the NetSim logs, not propagation latency alone. They include queueing and the interruptions that occur while satellite visibility and the serving beam change.

1.10 Conclusion

Over the 2.20-hour flight from Bangalore to Delhi, the serving satellite changes because elevation angle, slant range, beam geometry, path loss, and SINR all vary as the aircraft and the satellites move. Constellation size decides how often a usable satellite is overhead. With the current settings, the 100-satellite case delivers 45.50 Mbps of application throughput against 34.38 Mbps for the 50-satellite case, and the timelines show why: the denser constellation spends more of the flight connected.